

Attorney Docket No.
OI7011323001

ORACLE CONFIDENTIAL

UNITED STATES PATENT APPLICATION

FOR

**METHOD AND APPARATUS FOR MANAGING DATA ACCESS AND STORAGE OF
DATA LOCATED ON MULTIPLE STORAGE DEVICES**

INVENTORS:

WILLIAM BRIDGE

PREPARED BY:

PETER C. MEI

**BINGHAM MCCUTCHEN LLP
THREE EMBARCADERO CENTER, SUITE 1800
SAN FRANCISCO, CA 94111-4067**

ASSIGNEE: ORACLE INTERNATIONAL CORPORATION
500 ORACLE PARKWAY
REDWOOD SHORES, CA 94065

EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number EV348160044US

Date of Deposit March 15, 2004

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to MAIL STOP PATENT APPLICATION, Commissioner for Patents, Alexandria, VA 22313-1450.

Linda Major

(Typed or printed name of person mailing paper or fee)

(Signature of person mailing paper or fee)

SPECIFICATION**METHOD AND APPARATUS FOR MANAGING DATA ACCESS AND
STORAGE OF DATA LOCATED ON MULTIPLE STORAGE DEVICES**

5

BACKGROUND OF THE INVENTION**Field of the Invention**

The present invention relates to storage systems. More particularly, the invention relates to a method and apparatus for managing storage and access of data that is mirrored or copied to multiple storage locations.

10

Description of the Related Art

The present invention is directed to a method and system for improving I/O response times for applications that require optimal response times for data access, such as applications directed to storing or retrieving on-demand video or audio. Applications that provide video streams (or audio streams) often store portions of the video data in many locations on many data storage devices. As used herein, the term "data storage device" refers to any device or apparatus utilizable for the storage of data, e.g., a disk drive. For explanatory purposes only and not as an intent to limit the scope of this invention, the term "disk drive" will be used throughout this document instead of the term "data storage device."

20

To provide video on-demand at the proper display rate, the portions of stored video data are continually being retrieved and displayed at a viewing device. However,

different disk drives often provide different momentary response times to I/O requests due to varying system and environmental factors such as disk drive load, number of other requests on I/O request queues, network traffic, and physical disk drive proximity. If the video data is stored on many storage devices, then it is possible for the video data to be delivered out of order if some of the disk drives have momentary response times that are slower (or faster) than other disk drives.

One solution to this problem is to issue an I/O command for a particular piece of video data only when the prior sequential piece of video data has been delivered.

Alternatively, the system may buffer data so that even if the video stream arrives out of order, the video data is presented in order on the viewing device. The drawback to these solutions is that the delivered video stream is potentially stopped or constricted if any particular piece of video data has a relatively slow I/O response or delivery time. In other words, this method could produce an unacceptable display rate in the viewing device, resulting in a video stream that appears to have delayed or stopped motion on that display.

Another solution is to allow the video data to continue to stream even if the data arrives out of order, but to ignore any earlier pieces of data in the video sequence that arrive at the viewing device after a later piece of data in the video sequence. The drawback to this method is that gaps may appear in the video stream at the viewing device, resulting in a video stream having a jerky motion on the display.

Another solution is to store the video data (striped or unstriped) in an ordered manner on the disk drives. In this way, the chance of video data arriving out of order is reduced. One drawback to striping data in an ordered manner across disk drives is that it presents problems when new disk drives are added to the system. Adding new disk
5 drives may require wholesale restriping of data across the new configuration of disk drives. Moreover, storing data in an ordered fashion may cause load balancing problems, since the data cannot be randomly scattered across the disk drives.

Thus, there is a need for a method and system for allowing optimal response times
10 for access to data items stored in a storage system.

SUMMARY OF THE INVENTION

The invention encompasses a method and apparatus for managing storage and access of data to a storage system having a plurality of data storage devices.

One aspect of the invention is directed to reading a data item from a storage
5 system with optimized response time, in which multiple copies of the data items are stored to the storage system. The invention discloses a method of submitting more I/O read requests for the data item than are actually required to service the I/O needs of the node making the requests.

Another aspect of the invention is directed to writing and updating a data item,
10 whether mirrored or not, to a storage system. One feature of this aspect is the submission of more I/O write or update requests than are actually required to service the original I/O needs of the node making the requests.

Further details of aspects, objects, and advantages of the invention are described below in the detailed description, drawings, and claims. Both the foregoing general
15 description and the following detailed description are exemplary and explanatory, and are not intended to be limiting as to the scope of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings are included to provide a further understanding of the invention and, together with the Detailed Description, serve to explain the principles of the invention.

5

Fig. 1 illustrates an embodiment of a computer node accessing a disk storage system..

Fig. 2 illustrates an embodiment of a computer node having an I/O request queue accessing a disk storage system.

10

Fig. 3 is a flow diagram showing a process for reading data from a disk storage system.

Fig. 4 is a flow diagram showing a process for writing data to a disk storage system.

15

Fig. 5 is a flow diagram showing a process for updating data to a disk storage system.

Fig. 6 is a diagram of a computer hardware system with which the present invention can be implemented

DETAILED DESCRIPTION OF EMBODIMENT(S)

The present invention provides a method and system for increasing I/O response time by mirroring or replicating data at two or more locations, and submitting I/O requests to multiple storage locations for the same data. Different disk drives may
5 provide different response times to I/O requests due to varying system and environmental factors such as disk load, number of other requests on I/O request queues, network traffic, physical disk drive proximity, and unexpected failures. The present invention increases the likelihood of a faster response time to a particular request by sending the same request to two or more different disk drives that contain mirrored or replicated copies of
10 the same data. When a determination has been made as to the request that will produce the fastest response time, the other requests can thereafter be cancelled or ignored.

Mirroring involves the replication of data at two or more separate and distinct disk drives. To protect against the loss of information in the event of a disk drive failure,
15 mirroring can be employed to ensure continued access to data. However, mirroring can also be used to allow data retrieval from multiple storage locations. The present invention is applicable to any type of mirroring system, such as systems that mirror entire disk drives or systems that mirror only select partitions on a disk drive. Furthermore, the present invention is operable with systems that can mirror at granularities much smaller
20 than the entire disk drive, such as the storage system disclosed in co-pending U.S.

Application Ser. No. 09/177,916, entitled "Method And System For Managing Storage From Multiple Data Storage Devices", filed on October 23, 1998 and Application Ser. No. 09/178,387, entitled "Method and System for Managing Storage Systems Containing

Redundancy Data", filed on October 23, 1998, both of which are hereby incorporated by reference in their entirety.

To illustrate an embodiment of this invention, shown in Fig. 1 are disk drives 101
5 and 102 which are connected to disk controllers 110 and 112 respectively. Disk
controller 110 comprises a request queue for I/O requests made against disk drive 104
and disk controller 112 comprises a request queue 116 for I/O requests made against disk
drive 102. A particular data item is mirrored or replicated at extent 104 on disk 101 and
at extent 106 on disk 102.

10

If a client application on node 100 seeks to retrieve the data item, it may send an
I/O request to either disk 101 or 102 to read the contents of extents 104 or 106. However,
the momentary response time of one of the disk drives may be significantly slower than
the momentary response time of the other disk drive. This may occur, for example, if the
15 request queue for one of the disk drives is longer than the request queue for the other.
Thus, if node 100 queues its I/O request at random against one of the disk drives, there is
a chance that its request will be made against the disk drive having the slower response
time, and hence will result in a delayed response to the I/O request.

20 To address this problem, the method of the present invention comprises the
queuing of I/O requests for mirrored data at multiple data storage locations. According to
the invention, node 100 queues its I/O request to both disk drives 101 and 102. Thus, the
I/O request can be added to the request queue 114 on disk controller 110 and to the

request queue 116 on disk controller 112. Node 100 will then utilize the response from whichever disk controller/disk drive provides the fastest response time. Once node 100 receives a response from one disk controller/disk drive, a cancel command can be sent to the other disk controller/disk drive. The cancel command may take the form of an
5 instruction to remove the I/O request from the slower disk drive's corresponding request queue.

Alternatively, node 100 may allow both I/O operations to proceed to completion, even if a response has already been received from the faster disk drive, but will ignore
10 and/or not utilize the later response that arrives from the slower disk drive. Illustrated in Fig. 1 are I/O requests for mirrored data sent to only two disk drives, but the present invention is applicable to any number of disk drives containing mirrored or replicated data.

15 It is possible that more than one I/O request may complete at the same time. If multiple responses are received at node 100 at the same time, then only one is utilized, and the remaining responses are ignored/not utilized.

The present invention is applicable to systems having multiple levels of request
20 queues. Referring to the system illustrated in Fig. 2, shown is a computer node 200 having its own I/O request queue 202. Disk controllers 210 and 212 also have their own request queue to handle I/O requests for their corresponding disk drives 201 and 202. In the example of Fig. 2, an I/O request may be queued at either the node 200 itself (in

request queue 202), or at the disk controllers (in request queues 214 and 216). One reason for having multiple levels of queues is that the request queue for some disk controllers/disk drives permit only a single entry from each node on the system. Thus, the requesting node must maintain its own request queue to schedule and coordinate multiple I/O requests directed to that disk controller/disk drive.

When the invention is utilized in a system having multiple levels of request queues, I/O requests can be cancelled from any level. As an illustrative example, consider when a client application on node 200 seeks to access an item of data which is mirrored at disk locations 204 and 206 on disk drives 201 and 202 respectively. Node 200 may issue I/O requests to both disk drives, resulting in I/O requests that are initially added to request queue 202 on node 200. If request queues 214 and/or 216 have sufficient entry space, then the I/O requests can be removed from request queue 202 and sent to request queues 214 and/or 216. Once a response has been received for a first I/O request to one disk drive, the second I/O request to the other disk drive can be cancelled. Depending upon the configuration of the system, the second I/O request can be cancelled from either the request queue 202 on node 200, or if the I/O request has already been sent from request queue 202, it can be cancelled from the request queue for the second disk controller/disk drive.

20

Consistent with the invention, various events and conditions can be chosen to indicate that one disk drive will provide a faster response time than other disk drives, and thus will trigger cancellation or ignore commands for I/O requests sent to the other disk

drives having slower response times. In one embodiment, cancellation commands are issued to "slower" disk drives only when a first I/O request results in a fully delivered response from a disk drive (for a read request, this occurs when the requested data item has been fully retrieved and loaded into the requesting node's buffer). Another example of a possible triggering event is the movement of an I/O request from one request queue level to another. For example, if multiple I/O requests for a mirrored data item are loaded into request queue 202, the movement of a single one of those I/O requests to the request queue of a disk controller/disk drive could trigger the cancellation of the other related I/O requests still waiting in request queue 202. Depending upon the particular application to which the invention is directed, other triggering conditions and events can be employed within the scope of the invention.

The method of the present invention results in more requests that are placed on I/O queues than are actually required to service the immediate I/O needs of a client. This provides more even load balancing so that there are much more guarantees of service delivery. However, the present invention may result in multiple I/O requests sent to disk drives that either are not cancelled or cannot be cancelled. An uncancellable I/O request may occur, for example, if a particular disk controller does not allow external commands to remove an I/O request from its request queue. To address this, an embodiment of the system may be configured to allow multiple I/O requests to proceed to completion, but only the results from the fastest response disk drive is utilized. However, excess resource consumption may occur if the slower I/O requests cannot be cancelled, since some

systems may not permit resources allocated for uncancellable I/O requests to be released until all the related I/O requests have completed.

For example, the disk system may be obligated to allow all related I/O requests to
5 complete before reallocating memory in the node's buffer that is set aside for the
anticipated response to these requests. In an embodiment, this can be handled in two
ways. First, a single buffer allocation can be made for all related I/O requests for the
same mirrored data. Each response to the I/O requests writes to the same buffer location.
Therefore, regardless of how many I/O requests are queued, only a single buffer
10 allocation is made to hold the responses. Alternatively, a separate buffer allocation can
be made for each I/O request. When the first response is received at the requesting node,
the buffer allocation containing the first response is accessed and then freed. Thereafter,
each other buffer allocation is freed once their respective responses also arrive.

15 Significantly, the present invention does not require a common (or global) queue
to be operational. Thus, the present invention does not require a requesting node to be
able to contemporaneously view the contents of the request queue for any disk
controller/disk drive that it seeks to access. This is a key advantage over a related
technique to retrieve mirrored data from disk drives, which requires a node to be able to
20 view the request queues of the various disk drives containing the mirrored data and in
which the I/O request is sent only to the disk drive having the shortest queue. The
present invention provides a significant advantage over this technique since many disk
systems do not provide or allow external access to view disk drive request queues,

rendering the related technique inoperable. Moreover, even if an I/O request is sent to a particular disk drive having the numerically shortest request queue, other system and environmental factors (such as disk drive proximity and system workloads) may result in a slower response time from that disk drive as compared to other disk drives having longer request queues.

Fig. 3 illustrates the process flow for an embodiment of the invention that is used to perform read operations. The process for performing optimized reads according to the invention requires the storage system to maintain replicated or mirrored copies of the requested data items at two or more disk drives. As shown in Fig. 3, the following process actions are performed to read data according to an embodiment of the invention:

1. Submit I/O read operation requests to multiple disk drives that contain a mirrored or replicated copy of the requested data item (302). An I/O read operation request should be submitted to at least two disk drives that contain a copy of the data item. If more than two disk drives contain copies of the data item, then the requesting node may direct I/O requests to some or all of those disk drives.
2. Identify which disk drive is projected to provide the fastest response time to the I/O requests (304). In an embodiment, this determination is made by identifying the first disk drive to actually provide the requested data item to the requesting node. Other determination criteria may be implemented within the scope of the invention. For example, in an alternate embodiment, this

determination is made by identifying the first I/O read operation request to move from one level of request queues to another.

3. Test whether any of the other I/O read operation requests can be cancelled (306). This includes pending I/O read operation requests for the data item that have not yet been processed. If they exist, then cancellation commands (e.g., a command to remove the I/O request from a request queue) are issued to the identified I/O read operation requests that can be cancelled (308).

4. Test whether there exists other pending I/O read operation requests that are not or cannot be cancelled (310). I/O read operation requests may be uncancellable, if a particular disk controller/disk drive is not configured to allow external instructions to remove an I/O request from its request queue. The system may also choose not to incur the overhead of issuing cancellation commands for particular I/O read operations requests. If there exists any I/O requests that are not cancelled, then the requested data items may actually be retrieved from the respective disk drives responsive to the I/O read operations requests, but the requesting node will ignore and/or not utilize the data items that are retrieved (312).

The present invention can be employed to perform write operations to system data. Fig. 4 illustrates a process flowchart of the actions performed to execute the write operation in an embodiment of the invention:

1. Submit I/O write operation requests to multiple disk drives (402). The number of I/O requests should exceed the number of data copies that are intended

to be written out and saved to disks. For example, suppose that the data item is to be written out and saved in a triple mirrored system, i.e., there are intended to be three mirrored copies of the data item on three separate disk drives. Therefore, I/O write operations requests should be submitted to at least four or more disk drives to perform the desired write request. For a system in which mirroring is not used (i.e., only a single copy of the data item needs to be written), at least two or more I/O write operation requests should be submitted.

2. Identify which I/O write operation request is projected to provide the fastest response for the requested write operation (404). In an embodiment, this determination is made by identifying the first I/O write operation request that has actually completed the requested operation by writing the data item to disk. Other determination criteria may be implemented within the scope of the invention. For example, in an alternate embodiment, this determination is made by identifying the first I/O write operation request to move from one level of request queues to another.

3. Test whether a sufficient number of "fast" I/O requests has been identified (406). In a system that utilizes mirroring, more than one copy of a data item has to be written and saved to disk. For example, in a system having triple mirroring, the fastest three I/O requests are identified before discarding the other "slower" I/O write operation requests. Thus, if additional I/O request(s) need to be identified, the process returns back to process action 404 to identify an additional "fastest" I/O write operation requests. If a sufficient number of I/O write

operation requests have already been identified, then the process proceeds to process action 408.

4. Test whether there exists any pending I/O write operation requests for the data item that are to be cancelled (408). If so, then issue the appropriate cancellation commands (410). Otherwise, proceed to process act 412.

5. Test whether there exists any pending I/O write operation requests for the data item that are not to be cancelled or that cannot be cancelled (412). If they exist, then allow the I/O write operation requests to proceed to completion, but deallocate the disk space or delete the data items that are written to disk (414).

10

The present invention can be employed to perform update requests to system data. Update requests seek to modify or update data items that already exist on the system disk drives. Fig. 5 illustrates a process flowchart of the actions performed to execute update requests in an embodiment of the invention:

15 1. Submit I/O update operation requests to multiple disk drives (502). The number of I/O update operation requests should exceed the number of data copies that are intended to be maintained on the disk drives. For example, suppose that the data item is to be updated in a triple mirrored system, i.e., there are three mirrored copies of the data item on three separate disk drives that must be maintained. Therefore, the requesting node should submit I/O update operation requests to at least four or more disk drives to perform the update operation. For
20 a system in which mirroring is not used, I/O update operation requests should be submitted to at least two or more disk drives.

2. Identify which I/O update operation request is projected to provide the fastest response time for the requested update operation (504). In an embodiment, this determination is made by identifying the first I/O update operation request that has actually completed the requested update operation. Other determination
5 criteria may be employed within the scope of the invention. For example, in an alternate embodiment, this determination can be made by identifying the first I/O update operation request to move from one level of request queues to another.

3. In operation, the process of updating a data item is functionally the same as writing a data item to an unused section of a disk drive. Updating a data item
10 does not actually modify the existing data item already on disk, but writes the update data to a new and unused section of a disk drive. Once the updated data is written to a disk drive, pointers to the existing data are reconfigured to point to the newly written updated data (506). Further I/O read requests for the data item would be directed to this newly written section of the disk drive.

15 Alternatively, one or more of the I/O update operation requests may be directed to modifying the existing data item at an original disk drive location, rather than writing the updated data item to an unused section of a disk drive. If the "fastest" I/O update request is identified to be a request that updates an existing copy of the data item, then the pointers to this copy of the data item does
20 not have to be reconfigured.

In an embodiment, the pointers to the data item are not reconfigured until a sufficient number of "fast" I/O requests have been identified. At that point, pointers to existing copies of the data item that are not to be modified by one of

the identified "fast" I/O requests would be reconfigured in process action 506.

The existing copies of the data item which are no longer pointed to can be deleted from their respective disk drives. In some systems, this action may not be necessary once the original pointers to these sections of the disk drives have been redirected to the newly written sections of disk drives that contain the updated data.

4. Determine if more copies of the updated data need to be written to disk (508). If so, then return back to process action 504 and allow another pending I/O update operation request to proceed to completion. If only a single copy of the data item is to be written to disk (i.e., mirroring is not employed), then the first update actually written to a disk drive satisfies the required number of copies. However, if mirroring is employed, then additional copies of the updated data must be written to disk drives.

5. Once sufficient copies of the updated data have been written, test whether additional pending I/O update operation requests are to be cancelled (510). If true, then cancel the identified I/O update operation requests (512).

6. Test whether there exists pending I/O update operation requests for the same update that are not to be cancelled (514). This includes I/O requests which cannot be cancelled, as well as I/O requests that the system chooses not to cancel. If such I/O requests exist, then allow these requests to proceed to completion, but deallocate the disk drive space or delete the updated data that is written to disk as a result of these requests (516).

Hardware Overview

Figure 6 is a block diagram that illustrates a computer system 600 upon which an embodiment of the invention may be implemented. Computer system 600 includes a bus 602 or other communication mechanism for communicating information, and a processor 604 coupled with bus 602 for processing information. Computer system 600 also includes a main memory 606, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 602 for storing information and instructions to be executed by processor 604. Main memory 606 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 604. Computer system 600 further includes a read only memory (ROM) 608 or other static storage device coupled to bus 602 for storing static information and instructions for processor 604. A storage device 610, such as a magnetic disk or optical disk, is provided and coupled to bus 602 for storing information and instructions.

Computer system 600 may be coupled via bus 602 to a display 612, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 614, including alphanumeric and other keys, is coupled to bus 602 for communicating information and command selections to processor 604. Another type of user input device is cursor control 616, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 604 and for controlling cursor movement on display 612. This input device typically has two

degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

The invention is related to the use of computer system 600 for managing and allocating storage within a system comprising multiple data storage devices. According to one embodiment of the invention, such use is provided by computer system 600 in response to processor 604 executing one or more sequences of one or more instructions contained in main memory 606. Such instructions may be read into main memory 606 from another computer-readable medium, such as storage device 610. Execution of the sequences of instructions contained in main memory 606 causes processor 604 to perform the process steps described herein. One or more processors in a multi-processing arrangement may also be employed to execute the sequences of instructions contained in main memory 606. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

The term "computer-readable medium" as used herein refers to any medium that participates in providing instructions to processor 604 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 610. Volatile media includes dynamic memory, such as main memory 606. Transmission media includes coaxial cables, copper wire and fiber

optics, including the wires that comprise bus 602. Transmission media can also take the form of acoustic or light waves, such as those generated during radio wave and infrared data communications.

5 Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punch cards, paper tape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium
10 from which a computer can read.

 Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 604 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote
15 computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 600 can receive the data on the telephone line and use an infrared transmitter to convert the data to an infrared signal. An infrared detector coupled to bus 602 can receive the data carried in the infrared signal and place the data on bus 602. Bus 602
20 carries the data to main memory 606, from which processor 604 retrieves and executes the instructions. The instructions received by main memory 606 may optionally be stored on storage device 610 either before or after execution by processor 604.

Computer system 600 also includes a communication interface 618 coupled to bus 602. Communication interface 618 provides a two-way data communication coupling to a network link 620 that is connected to a local network 622. For example, communication interface 618 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 618 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 618 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

Network link 620 typically provides data communication through one or more networks to other data devices. For example, network link 620 may provide a connection through local network 622 to a host computer 624 or to data equipment operated by an Internet Service Provider (ISP) 626. ISP 626 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 628. Local network 622 and Internet 628 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 620 and through communication interface 618, which carry the digital data to and from computer system 600, are exemplary forms of carrier waves transporting the information.

Computer system 600 can send messages and receive data, including program code, through the network(s), network link 620 and communication interface 618. In the Internet example, a server 630 might transmit a requested code for an application program through Internet 628, ISP 626, local network 622 and communication interface 618. In accordance with the invention, one such downloaded application provides for managing, storing, and retrieving data from a storage system containing multiple data storage devices. The received code may be executed by processor 604 as it is received, and/or stored in storage device 610, or other non-volatile storage for later execution. In this manner, computer system 600 may obtain application code in the form of a carrier wave.

In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention. For example, the above-described process flows are described with reference to a particular ordering of process actions. However, the ordering of many of the described process actions may be changed without affecting the scope or operation of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than restrictive sense.